arXiv:1610.09263v1 [cs.AI] 28 Oct 2016

Flexible constrained sampling with guarantees for pattern mining

Vladimir Dzyuba[1], Matthijs van Leeuwen[2], Luc De Raedt[1]

Abstract

Pattern sampling has been proposed as a potential solution to the infamous pattern explosion. Instead of enumerating all patterns that satisfy the constraints, individual patterns are sampled proportional to a given quality measure. Several sampling algorithms have been proposed, but each of them has its limitations when it comes to 1) flexibility in terms of quality measures and constraints that can be used, and/or 2) guarantees with respect to sampling accuracy.

We therefore present Flexics, the first flexible pattern sampler that supports a broad class of quality measures and constraints, while providing strong guarantees regarding sampling accuracy. To achieve this, we leverage the perspective on pattern mining as a constraint satisfaction problem and build upon the latest advances in sampling solutions in SAT as well as existing pattern mining algorithms. Furthermore, the proposed algorithm is applicable to a variety of pattern languages, which allows us to introduce and tackle the novel task of sampling sets of patterns.

We introduce and empirically evaluate two variants of Flexics: 1) a generic variant that addresses the well-known itemset sampling task and the novel pattern set sampling task as well as a wide range of expressive constraints within these tasks, and 2) a specialized variant that exploits existing frequent itemset techniques to achieve substantial speed-ups. Experiments show that Flexics is both accurate and efficient, making it a useful tool for pattern-based data exploration.

1 Introduction

Pattern mining [1] is an important and well-studied task in data mining. Informally, a pattern is a statement in a formal language that concisely describes a subset of a given dataset. Pattern mining techniques aim at providing comprehensible descriptions of coherent regions in the data. Many variations of pattern

[1] DTAI, KU Leuven, Belgium, firstname.lastname@cs.kuleuven.be
[2] LIACS, Leiden University, The Netherlands, m.van.leeuwen@liacs.leidenuniv.nl

mining have been proposed in the literature, together with even more algorithms to efficiently mine the corresponding patterns. Best known is frequent pattern mining [2], which includes frequent itemset mining and its extensions.

Traditional pattern mining methods enumerate all frequent patterns, though it is well-known that this usually results in humongous amounts of patterns (the infamous *pattern explosion*). To make pattern mining more useful for exploratory purposes, different solutions to this problem have been proposed. Each of these solutions has its own advantages and disadvantages. *Condensed representations* [3] can often be efficiently mined, but generally still result in large numbers of patterns. *Top-k mining* [4] is efficient but results in strongly related, redundant patterns showing a lack of diversity. *Constrained mining* [5] may result in too few or too many patterns, depending on the user-chosen constraints. *Pattern set mining* [6] takes into account the relationships between the patterns, which can result in small solution sets, but is computationally intensive.

In this paper, we study *pattern sampling*, another approach that has been proposed recently: instead of enumerating all patterns, patterns are sampled one by one, according to a probability distribution that is proportional to a given quality measure. The promised benefits include: 1) flexibility in that potentially a broad range of quality measures and constraints can be used; 2) 'anytime' data exploration, where a growing representative set of patterns can be generated and inspected at any time; 3) diversity in that the generated sets of patterns are independently sampled from different regions in the solution space. To be reliable, pattern samplers should provide theoretical guarantees regarding the sampling accuracy, i.e., the difference between the empirical probability of sampling a pattern and the (generally unknown) target probability determined by its quality.

While a number of pattern sampling approaches has been developed over the past years, they are either inflexible (as they only support a limited number of quality measures and constraints), or do not provide theoretical guarantees concerning the sampling accuracy. At the algorithmic level, they follow standard sampling approaches such as Markov Chain Monte-Carlo random walks over the pattern lattice [7, 8, 9], or a special purpose sampling procedure tailored for a restricted set of itemset mining tasks [10, 11]. Although MCMC approaches are in principle applicable to a broad range of tasks, they often converge only slowly to the desired target distribution and require the selection of the "right" proposal distributions.

To the best of our knowledge, none of the existing approaches to pattern sampling takes advantage of the latest developments in sampling technology from the SAT-solving community, where a number of powerful samplers have been developed based on random hash functions and XOR-sampling [12, 13, 14, 15]. WeightGen [16], one of the recent approaches, possesses the benefits mentioned above: it is an any-time algorithm, it is flexible as it works with any distribution, it generates diverse solutions, and provides strong performance guarantees under reasonable assumptions.

In this paper, we show that the latest developments in sampling solutions

Table 1: Our method is the first pattern sampler that combines flexibility with respect to the choice of constraints and sampling distributions with strong theoretical guarantees.

Sampler	Arbitrary constraints	Arbitrary distributions	Strong guarantees	Efficiency	Pattern set sampling
ACFI [7]	Minimal frequency	-	-	✓	-
LRW [8]	✓	✓	-	-	-
FCA [9]	Anti-/ monotonic	✓	-	✓	-
TS (Two-step) [10, 11]	-	-	✓	✓	-
FLEXICS This paper	GFLEXICS	✓	✓	EFLEXICS	✓

in SAT are also relevant to pattern sampling and essentially offer the same advantages. Our results build upon the view of pattern mining as constraint satisfaction, which is widely accepted in the data mining community [17].

Approach and contributions More specifically, we introduce FLEXICS: a flexible pattern sampler that samples from distributions induced by a variety of pattern quality measures and allows for a broad range of constraints while still providing strong theoretical guarantees. Notably, FLEXICS is, in principle, agnostic of the quality measure, as the sampler treats it as a black box. (However, note that its properties affect the efficiency of the algorithm.) The other building block is a *constraint oracle* that enumerates all patterns that satisfy the constraints, i.e., a mining algorithm. The proposed approach allows converting an existing pattern mining algorithm into a sampler with guarantees. Thus, its flexibility is not limited by the choice of constraints and quality measures, but even allows tackling richer pattern languages, which we demonstrate by tackling the novel task of *sampling sets of patterns*. Table 1 compares the proposed approach to alternative samplers; see Section 3 for a more detailed discussion.

The main technical contribution of this paper consists of two versions of the FLEXICS sampler, which are based on different constraint oracles. First, we introduce a generic variant, dubbed GFLEXICS, that supports a wide range of pattern constraints, such as syntactic or redundancy-eliminating constraints. GFLEXICS uses CP4IM [17], a declarative constraint programming-based mining system, as its oracle. Any constraint supported by CP4IM can be used without interfering with the umbrella procedure that performs the actual sampling task. We use GFLEXICS to tackle a wide range of well-known itemset sampling tasks

as well as the novel pattern set sampling task. Second, as it is well-known that generic solvers impose an overhead on runtime, we introduce a variant specialized towards frequent itemsets, dubbed EFlexics, which has an extended version of Eclat [18] at its core as oracle.

Experiments show that Flexics' sampling accuracy is impressively high: in a variety of settings supported by the sampler, empirical frequencies are within a small factor of the target distribution induced by various quality measures. Furthermore, practical accuracy is substantially higher than theory guarantees. EFlexics is shown to be faster than its generic cousin, demonstrating that developing specialized solvers for specific tasks is beneficial when runtime is an issue. Finally, the flexibility of the sampler allows us to use the same approach to successfully tackle the novel problem of sampling pattern sets. This demonstrates that Flexics is a useful tool for pattern-based data exploration.

This paper is organized as follows. We formally define the problem of pattern sampling in Section 2. After reviewing the related research in Section 3, we present the two key ingredients of the proposed approach in Section 4: 1) the perspective on pattern mining as constraint satisfaction and 2) hashing-based sampling with WeightGen. In Section 5, we present Flexics, a flexible pattern sampler with guarantees. In particular, we detail modifications required to adapt WeightGen to pattern sampling and describe the procedure to convert two existing mining algorithms into oracles suitable for use with WeightGen, which yields two variants of Flexics. In Section 6, we introduce the pattern set sampling task and describe how it can be tackled with Flexics. Furthermore, we outline sampling non-overlapping tilings, an example of pattern set sampling that is studied in the experiments. The experimental evaluation in Section 7 investigates the accuracy, scalability, and flexibility of the proposed sampler. We present our conclusions in Section 8.

2 Problem definition

Here we present a high-level definition of the task that we consider in this paper; for concrete instances and examples, see Sections 4 and 6. The pattern sampling problem is formally defined as follows: given a dataset $\mathcal{D}$, a pattern language $\mathcal{L}$, a set of constraints $\mathcal{C}$, and a quality measure $\varphi : \mathcal{L} \to \mathbb{R}_+$, generate random patterns that satisfy constraints in $\mathcal{C}$ with probability proportional to their qualities:

$$P_\varphi(p) = \begin{cases} \varphi(p)/Z_\varphi & \text{if } p \in \mathcal{L} \text{ satisfies } \mathcal{C} \\ 0 & \text{otherwise} \end{cases}$$

where Z_φ is an (unknown) normalization constant.

A quality measure quantifies the domain-specific interestingness of a pattern. The choice of a quality measure and constraints allows a user to express her analysis requirements. The sampling procedure meets these requirements by satisfying the constraints and generating high-quality patterns more frequently.

Thus, sampled patterns are a representative subset of all interesting regularities in the dataset.

Pattern set mining is an extension of pattern mining, which considers sets of patterns rather than individual patterns. Despite its popularity, we are not aware of the existence of pattern set samplers. The task of pattern set sampling can easily be formalized as an extension of pattern sampling, where we sample sets of patterns $s \subset \mathcal{L}$, and the constraints $\mathcal{C}$ as well as the quality measure φ are specified over sets of patterns (from $2^{\mathcal{L}}$) rather than individual patterns (from $\mathcal{L}$).

3 Related work

We here focus on two classes of related work, i.e., 1) pattern mining as constraint satisfaction and 2) pattern sampling.

Constrained pattern mining The study of constraints has been a prominent subfield of pattern mining. A wide range of constraint classes were investigated, including anti-monotonic constraints [1], convertible constraints [19], and others. Another development of these ideas led to the introduction of global constraints that concern multiple patterns and to the emergence of *pattern set mining* [20, 21]. Furtheremore, generic mining systems that could freely combine various constraints were proposed [22, 23].

These insights allowed to draw a connection between pattern mining and constraint satisfaction in AI, e.g., SAT or constraint programming (CP). As a result, declarative mining systems, which use generic constraint solvers to mine patterns according to a declarative specification of the mining task, were proposed. For example, CP was used to develop first declarative systems for itemset mining [17] and pattern set mining [24, 25]. Recently, declarative approaches have been extended to support sequence mining [26] and graph mining [27].

Constraint-based systems allow a user to specify a wide range of pattern constraints and thus provide tools to alleviate the pattern explosion. However, the underlying solvers use systematic search, which affects the order in which the patterns are generated and hence the diversity. Similarly, pattern set miners that directly aim at obtaining diverse result sets typically incur prohibitive computational costs as the size of the pattern space grows.

Pattern sampling Table 1 compares our method with the approaches described in Section 1, namely MCMC and *two-step* samplers [10, 11]. We further break down MCMC samplers into three groups: *ACFI*, the very first uniform sampler developed for approximate counting of frequent itemsets [7]; *LRW*, a generic approach based on random walks over pattern lattice [8]; and *FCA*, a sampler, which uses Markov chains based on insights from formal concept analysis [9].

Although MCMC samplers provide theoretical guarantees, in practice, their convergence is often slow and hard-to-diagnose. Such solutions as long burn-in or heuristic adaptations either increase the runtime or weaken the guarantees. Furthermore, ACFI is tailored for a single task; FCA only supports anti-/monotone constraints; and LRW checks constraints locally, while building

the neighborhood of a state, which might require advanced reasoning. Two-step samplers, while provably accurate and efficient, only support a limited number of weight functions and do not support constraints.

4 Preliminaries

We first introduce itemset mining, a prototypical pattern mining task, and formalize it as a CSP and then describe WeightGen, a hashing-based sampling algorithm.

4.1 Itemset mining

Itemset mining is an instance of pattern mining specialized for binary data. Let $\mathcal{I} = \{1 \dots M\}$ denote a set of items. A dataset $\mathcal{D}$ is a bag of transactions over $\mathcal{I}$, where each transaction t is a subset of $\mathcal{I}$, i.e., $t \subseteq \mathcal{I}$; $\mathcal{T} = \{1 \dots N\}$ is a set of transaction indices. The pattern language $\mathcal{L}$ also consists of sets of items, i.e., $\mathcal{L} = 2^{\mathcal{I}}$. An itemset p occurs in a transaction t, iff $p \subseteq t$. The frequency of p is the number of transactions in which it occurs: $freq\,(p) = |\{t \in \mathcal{D} \mid p \subseteq t\}|$. In labeled datasets, a transaction has a label from $\{-, +\}$; $freq^{-,+}$ are defined accordingly.

We first give a brief overview of the general approach to solving CSPs and then present a formalization of itemset mining as a CSP, following that of CP4IM [17]. Formally, a CSP is comprised of *variables* along with their *domains* and *constraints* over these variables. The goal is to find a solution, i.e., an assignment of values to all variables that satisfies all constraints. Every constraint is implemented by a *propagator*, i.e., an algorithm that takes domains as input and removes values that do not satisfy the constraint. Propagators are activated when variable domains change, e.g., by the search mechanism or other propagators. A CSP solver is typically based on depth-first search. After a variable is assigned a value, propagators are run until domains cannot be reduced any further. At this point, three cases are possible: 1) a variable has an empty domain, i.e., the current search branch has failed and backtracking is necessary, 2) there are unassigned variables, i.e., further branching is necessary, or 3) all variables are assigned a value, i.e., a solution is found.

Let I_i denote a variable corresponding to each item; T_t a variable corresponding to each transaction; and $\mathcal{D}_{ti}$ a constant that is equal to 1, if item i occurs in transaction t, and 0 otherwise. Variables I_i and T_t are *binary*, i.e., their domain is $\{0, 1\}$. Each CSP solution corresponds to a single itemset. Thus, for example, $I_i = 1$ implies that item i is included in the current (partial) solution, whereas $T_t = 0$ implies that transaction t is *not* covered by it. Table 2 lists some of the most common constraints. The *coverage* constraint essentially models a dataset query and ensures that if the item variable assignment corresponds to an itemset p, only those transaction variables that correspond to indices of transactions where p occurs, are assigned value 1. Other constraints allow users to remove uninteresting solutions, e.g., redundant non-*closed* itemsets. Most solvers pro-

Table 2: Constraint programming formulations of common itemset mining constraints. $I_i = 1$ implies that item i is included in the current (partial) solution, whereas $T_t = 1$ implies that it covers transaction t.

Constraint	Parameters	CP formulation
$coverage$		$\forall t \in \mathcal{T} \;\; T_t = 1 \;\; \Leftrightarrow \;\; \sum_{i \in \mathcal{I}} I_i \left(1 - \mathcal{D}_{ti}\right) = 0$
$minfreq\left(\theta\right)$	$\theta \in (0,\ 1]$	$\forall i \in \mathcal{I} \;\; I_i = 1 \;\; \Rightarrow \;\; \sum_{t \in \mathcal{T}} T_t \mathcal{D}_{ti} \geq \theta \times \left\lvert\mathcal{D}\right\rvert$
$closed$		$\forall i \in \mathcal{I} \;\; I_i = 1 \;\; \Leftrightarrow \;\; \sum_{t \in \mathcal{T}} T_t \left(1 - \mathcal{D}_{ti}\right) = 0$
$minlen\left(\lambda\right)$	$\lambda \in [1,\ M]$	$\forall t \in \mathcal{T} \;\; T_t = 1 \;\; \Rightarrow \;\; \sum_{i \in \mathcal{I}} I_i \mathcal{D}_{ti} \geq \lambda$

vide facilities for enumerating all solutions in sequence, i.e., to enumerate all patterns.

In contrast to hard constraints, *quality measures* are used to describe soft user preferences with respect to interestingness of patterns. Common quality measures concern frequency, e.g., $\varphi \equiv freq$, discriminativity in a labeled dataset, e.g., *purity* $\varphi\left(p\right) = \max\left\{freq^{+}\left(p\right), freq^{-}\left(p\right)\right\} / freq\left(p\right)$, etc.

4.2 WeightGen

WeightGen [16] is an algorithm for approximate weighted sampling of satisfying assignments (solutions) of a Boolean formula that only requires access to an efficient constraint oracle that enumerates the solutions, e.g., a SAT solver. The core idea consists in partitioning the solution space into a number of "cells" and sampling a solution from a random cell. Partitioning with desired properties is obtained via augmenting the original problem with random XOR constraints. Theoretical guarantees stem from independence properties of uniformly generated XOR constraints. The sequel follows Sections 3-4 in [16].

Problem statement and guarantees Formally, let $\mathbf{F}$ denote a Boolean formula; F a satisfying variable assignment of $\mathbf{F}$; M the total number of variables; $w\left(\cdot\right)$ a black-box weight function that for each F returns a number in $(0,\ 1]$; and w_{min} (resp. w_{max}) the minimal (resp. maximal) weight over all satisfying assignments of $\mathbf{F}$. The weight function induces the probability distribution over satisfying assignments of $\mathbf{F}$, where $P_w\left(F\right) = w\left(F\right) / \sum w\left(F'\right)$. Quantity $r = w_{max} / w_{min}$ is the (possibly unknown) *tilt* of the distribution P_w.

Given a user-provided upper bound on tilt $\hat{r} \geq r$ and a desired sampling error tolerance $\kappa \in (0,\ 1)$ (the lower κ, the tighter the bounds on the sampling error), WeightGen generates a random solution F. Performance guarantees concern both accuracy and efficiency of the algorithm and depend on the parameters and the number of variables M; see Section 5 for details.

Algorithm Recall that the core idea that underlies sampling with guarantees is partitioning the overall solution space into a number of random cells by adding random XOR constraints. WeightGen proceeds in two phases: 1) the estimation phase and 2) the sampling phase. The goal of the *estimation phase*

is to estimate the number of XOR constraints necessary to obtain a "small" cell, where the required cell weight is determined by the desired sampling error tolerance.

The *sampling phase* starts with applying the estimated number of XOR constraints. If it obtains a cell whose total weight lies within a certain range that depends on κ, a solution is sampled exactly from all solutions in the cell; otherwise, it adds a new random XOR constraint. However, the number of XOR constraints that can be added is limited. If the algorithm cannot obtain a suitable cell, it indicates failure and returns no sample.

Both phases make use of a *bounded* oracle that terminates as soon as the total weight of enumerated solutions exceeds a predefined number. It enumerates solutions of the original problem $\mathbf{F}$ augmented with the XOR constraints. An individual XOR constraint over variables $\mathbf{X}$ has the form $\bigotimes b_i \cdot X_i = b_0$, where $b_{0|i} \in \{0, 1\}$. The coefficients b_i determine the variables involved in the constraint, whereas the *parity bit* b_0 determines whether an even or an odd number of variables msut be set to 1. Together, m XOR constraints identify one cell belonging to a partitioning of the overall solution space into 2^m cells.

The core operation of WeightGen involves drawing coefficients uniformly at random, which induces a random partitioning of the solution space that satisfies the 3-*wise independence property*, i.e., knowing the cells for two arbitrary assignments does not provide any information about the cell for a third assignment [12]. This ensures desired statistical properties of random partitions, required for the theoretical guarantees. The reader interested in further technical details should consult Appendix A and [16].

5 Flexics: Flexible pattern sampler with guarantees

In this paper, we propose Flexics, a pattern sampler that uses WeightGen as the umbrella sampling procedure. To this end, we 1) extend it to CSPs with binary variables, a class of problems that is more general than SAT and that includes pattern mining as described in Section 4; 2) augment existing mining algorithms for use with WeightGen; and 3) investigate the properties of pattern quality measures in the context of WeightGen's requirements.

WeightGen was originally presented as an algorithm to sample solutions of the SAT problem. Pattern mining problems cannot be efficiently tackled by pure Boolean solvers due to the prominence of cardinality constraints (e.g., $minfreq$). However, we observe that the core sampling procedure is applicable to any CSP with binary variables, as its solution space can be partitioned with XOR constraints in the required manner.

Based on this insight, we present two variants of Flexics that differ in their oracles. Each oracle is essentially a pattern mining algorithm extended to support XOR constraints along with common constraints on patterns. The first one, dubbed GFlexics, builds upon the generic formalization and solving

techniques described in Section 4 and thus supports a wide range of constraints. Owing to the properties of the *coverage* constraint, XOR constraints only need to involve item variables[1], which makes them relatively short, mitigating the computational overhead. Moreover, this perspective helps us design the second approach, dubbed EFlexics, which uses an extension of Eclat [18], a well-known mining algorithm, as an oracle. It is tailored for a single task (frequent itemset mining, i.e., it only supports the $minfreq$ constraint), but is capable of handling larger datasets. We describe each oracle in detail in the following subsections.

Given a dataset $\mathcal{D}$, constraints $\mathcal{C}$, a quality measure φ, and the error tolerance parameter $\kappa \in (0,\ 1)$, Flexics first constructs a CSP corresponding to the task of mining patterns satisfying $\mathcal{C}$ from $\mathcal{D}$. It then determines parameters for the sampling procedure, including the appropriate number of XOR constraints, and starts generating samples. To this end, it uses one of the two proposed oracles to enumerate patterns that satisfy $\mathcal{C}$ *and* random XOR constraints. Both variants of Flexics support sampling from black-box distributions derived from quality measures and, most importantly, preserve the theoretical guarantees of WeightGen[2]:

Theorem 1. *The probability that* Flexics *samples a random pattern p that satisfies constraints $\mathcal{C}$ from a dataset $\mathcal{D}$, lies within a bounded range determined by the quality of the pattern $\varphi(p)$ and κ:*

$$\frac{\varphi(p)}{Z_\varphi} \times \frac{1}{1+\varepsilon(\kappa)} \leq P(\textsc{Flexics}(\mathcal{D},\mathcal{C},\varphi;\kappa)=p) \leq \frac{\varphi(p)}{Z_\varphi} \times (1+\varepsilon(\kappa))$$

Proof. Theorem 3 of [16] states:

$$P_w(F)/(1+\varepsilon(\kappa)) \leq \hat{P}_F \leq P_w(F) \times (1+\varepsilon(\kappa))$$

where $\hat{P}_F$ denotes the probability that WeightGen called with parameters $\hat{r}$ and κ samples the solution F, $P_w(F) \propto w(F)$ denotes the target probability of F, and $\varepsilon(\kappa) = (1+\kappa)\left(2.36 + 0.51/(1-\kappa)^2\right) - 1$ denotes sampling error derived from κ.

For technical purposes, we introduce the notion of the *weight* of a pattern as its quality scaled to the range $(0,\ 1]$, i.e., $w_\varphi(p) = \varphi(p)/C$, where C is an arbitrary constant such that $C \geq \max_{p\in\mathcal{L}} \varphi(p)$. The proof follows from Theorem 3 of [16] and the observation that $\textsc{Flexics}(\mathcal{D},\mathcal{C},\varphi;\kappa)$ is equivalent to $\textsc{WeightGen}(\text{CSP}(\mathcal{D},\mathcal{C}), w_\varphi;\kappa)$. The estimation phase effectively corrects potential discrepancy between C and Z_φ. □ □

Furthemore, Theorem 4 by [16], provides *efficiency guarantees*: the number of calls to the oracle is linear in $\hat{r}$ and polynomial in M and $1/\varepsilon(\kappa)$. The

[1] In other words, item variables I are the *independent support* of a pattern mining CSP.
[2] Theorem 1 corresponds to and follows from Theorem 3 of [16].

assumption that the tilt is *bounded from above* by a reasonably low number is the only assumption regarding a (black-box) weight function. Moreover, it only affects the efficiency of the algorithm, but not its accuracy.

Thus, using a quality measure with FLEXICS requires knowledge of two properties: scaling constant C and tilt bound $\hat{r}$. In practice, both are fairly easy to come up with for a variety of measures. For example, for $freq$ and *purity*, $C = |\mathcal{D}|$, $\hat{r} = \theta^{-1}$ and $C = 1$, $\hat{r} = 2$ respectively; see Section 6 for another example.

5.1 GFlexics: Generic pattern sampler

The first variant relies on CP4IM, a constraint programming-based mining system. A wide range of constraints supported by CP4IM are automatically supported by the sampler and can be freely combined with various quality measures.

In order to turn CP4IM into a suitable bounded oracle, we need to extend it with an efficient propagator for XOR constraints. It is based on the process of *Gaussian elimination* [28], a classical algorithm for solving systems of linear equations. Each XOR constraint can be viewed as a linear equality over the field $\mathbb{F}_2$ of two elements, 0 and 1, and all coefficients form a binary matrix (Figure 1.2). At each step, the matrix is updated with the latest variable assignments and transformed to *row echelon form*, where all ones are on or above the main diagonal and all non-zero rows are above any rows of all zeroes (Figure 1.3). During echelonization, two situations enable propagation. If a row becomes empty while its right hand side is equal to 1, the system is unsatisfiable and the current search branch terminates (Figure 1.5). If a row contains only one free variable, it is assigned the right hand side of the row (Figure 1.3).

Gaussian elimination in $\mathbb{F}_2$ can be performed very efficiently, because no division is necessary (all coefficients are 1), and subtraction and addition are equivalent operations. For a system of k XOR constraints over n variables, the total time complexity of Gaussian elimination is $\mathcal{O}\left(k^2 n\right)$.

				↓ ↓
$x_1 \otimes x_5 = 1$	1 0 0 0 1 | 1	1 0 0 0 1 | 1	1 0 0 0 1 | 1	→ **0** 0 0 0 **0** | **1**
$x_2 \otimes x_3 \otimes x_4 \otimes x_5 = 0$	0 1 1 1 1 | 0	→ 0 **1** 0 0 0 | **0**	0 0 0 1 1 | 1	**0** 0 0 1 **0** | **0**
$x_1 \otimes x_2 \otimes x_3 \otimes x_5 = 0$	1 1 1 0 1 | 0	→ 0 0 **1** 0 0 | **1**	0 0 0 0 0 | 0	0 0 0 0 0 | 0
$x_2 \otimes x_4 \otimes x_5 = 1$	0 1 0 1 1 | 1	0 0 0 1 1 | 1	0 0 0 0 0 | 0	0 0 0 0 0 | 0
1) Random XOR constraints	2) Initial constraint matrix	3) Echelonized matrix: x_2 and x_3 are assigned a value	4) Updated matrix	5) If x_1 and x_5 are set to 1 (e.g., by search), the system is unsatisfiable

Figure 1: Propagating XOR constraints using Gaussian elimination in $\mathbb{F}_2$.

5.2 EFlexics: Efficient pattern sampler

Generic constraint solvers currently cannot compete with efficiency and scalability of specialized mining algorithms. In order to develop a less flexible, yet more efficient version of our sampler, we extend the well-known Eclat algorithm to handle XOR constraints. Thus, EFlexics is tailored for frequent itemset sampling and uses EclatXOR (Algorithm 1) as an oracle.

Algorithm 1 shows the pseudocode of the extended Eclat. The algorithm relies on the *vertical* data representation, i.e., for each candidate item, it stores a set of indices of transactions (TIDs), in which this item occurs (Line 4). Eclat starts with determining frequent items and ordering them, by frequency ascending. It explores the search space in a depth-first manner, where each branch corresponds to (ordered) itemsets that share a prefix.

The core operation is referred to as *processing an equivalence class* of itemsets (EqClass). For each prefix, Eclat maintains a set of candidate suffixes, i.e., items that follow the last item of the prefix in the item order and are frequent. The frequency of a candidate suffix, given the prefix, is computed by intersecting its TID with the TID of the prefix (Lines 9, 15, and 22).

We extend Eclat with XOR constraint handling (Lines 16-22). Variable updates stem from Eclat extending the prefix and removing infrequent suffixes (Line 16). XOR propagation can result in extending the prefix or removing candidate suffixes as well (Line 19). Furthermore, if the prefix has been extended, TIDs of candidate suffixes need to be updated, with some of them possibly becoming infrequent, leading to further propagation (Lines 19-22). If the prefix becomes infrequent, the search branch terminates.

Fixed variable-order search, like Eclat, is an advantageous case for Gaussian elimination [29]: non-zero elements are restricted to the right region of the matrix, hence Gaussian elimination only needs to consider a contiguous, progressively shrinking subset of columns. Total memory overhead of EclatXOR compared to plain Eclat is $\mathcal{O}\left(d \times |\mathcal{F}| \times N_{XOR} + pivot \times r\right)$, where d denotes maximal search depth, $|\mathcal{F}|$ the number of frequent singletons (columns of a matrix), and N_{XOR} the number of XOR constraints (rows of a matrix). The first term refers to a set of XOR matrices in unexplored search branches, whereas the second term refers to storing itemsets in a cell (Line 19 in Algorithm 2 in Appendix A).

6 Pattern set sampling

We highlight the flexibility of Flexics by introducing and tackling the novel task of *sampling sets of patterns*. For the purposes of sampling, a set of patterns is essentially treated as a composite pattern. Typically, constituent patterns are required to be different from each other. The quality (and hence, the sampling probability) of a pattern set depends on collective properties of constituent patterns. These characteristics, coupled with the immense size of the pattern set search space, make sampling even more challenging.

Algorithm 1 ECLAT augmented with XOR constraint propagation (Lines 16-22)

Input: Dataset $\mathcal{D}$ over items $\mathcal{I}$, min.freq θ, XOR matrix M
Assumes: Item order $\succ_{\mathcal{I}}$ by frequency ascending

1: **function** ECLATXOR($\mathcal{D}$, θ, M)
▷ Mine all frequent patterns that satisfy XOR constraints encoded by M
2: Frequent items $FI = \emptyset$
3: **for** item $i \in \mathcal{I}$ **do**
4: $TID_i = \{\text{transaction index } t \in \mathcal{T} \mid \mathcal{D}_{ti} = 1\}$
5: **if** $|TID_i| \geq \theta$ **then** ▷ Item is frequent
6: $FI \overset{Add}{\leftarrow} (i, TID_i)$
7: SORT(FI, $\succ_{\mathcal{I}}$)
8: **for** $i \in FI$ **do**
9: Candidate suffixes $CS = \{i' \in FI \setminus i \mid i' >_{\mathcal{I}} i\}$
10: EQCLASS($\{i\}$, CS, M)
11: **function** EQCLASS(Prefix P, cand.suffixes $CS \neq \emptyset$, M)
▷ Mine all patterns that start with P
12: **if** CHECKCONSTRAINTS(P, M) **then**
13: **return** P ▷ Return prefix, if it satisfies XORs
14: **for** candidate suffix $s \in CS$ **do**
15: $P' = P \cup s$; frequent suffixes $FS = \{f \in CS \setminus s \mid f >_{\mathcal{I}} s \wedge |f.TID \cap s.TID| \geq \theta\}$

16: $U_1 = \{s\}$, $U_0 = CS \setminus FS$ ▷ Variable updates
17: $M' =$ UPDATEANDECHELONIZE(M, U_1, U_0)
18: $(A_1, A_0) =$ PROPAGATE(M') ▷ Item variables
▷ that were assigned value 1 or 0 by propagation
19: $FS' = FS \setminus (A_1 \cup A_0)$
20: **if** $A_1 \neq \emptyset$ **then** ▷ If prefix was extended,
▷ update TIDs and check support
21: $P' \leftarrow P' \cup A_1$, $\Delta_{TID} = \bigcap_{f \in A_1} f.TID$
22: $FS' \leftarrow \{f' \in FS' : |f'.TID \cap \Delta_{TID}| \geq \theta\}$

23: **if** $|P'.TID| \geq \theta \wedge FS' \neq \emptyset$ **then**
24: EQCLASS(P', FS'', M')

To develop a sampler, we extend GFLEXICS with the CSP-formulation of the k-pattern set mining task [25], which in turn builds upon the formulation of the itemset mining task described in Section 4. Recall that a CSP is defined by a set of variables and constraints over these variables. Each constituent pattern is modeled with distinct item and transaction variables, item, item variables I_{ik} and T_{tk} for the kth pattern p_k. Note that this increases the length of XOR constraints, which poses an additional challenge from the sampling perspective.

Any single-pattern constraint can be enforced for a constituent pattern, e.g., $minfreq\,(\theta)$, *closed*, or $minlen\,(\lambda)$. A common pattern set-specific constraint is *no overlap*, which enforces that neither the itemsets (1), nor the sets of transactions that they cover (2) overlap:

$$(1)\ \forall i \in \mathcal{I}\ \sum I_{ik} \leq 1 \qquad\qquad (2)\ \forall t \in \mathcal{T}\ \sum T_{tk} \leq 1$$

Furthermore, there is typically a symmetry-breaking constraint that requires that the set of transaction indices of p_i lexicographically precedes those of $\{p_j,\ j > i\}$. This approach allows modeling a wide range of pattern set sampling tasks, e.g., sampling k-term DNFs, conceptual clusterings, redescriptions, and others. In this paper, we use the problem of *tiling datasets* [30] as an example.

The main aim of tiling is to cover a large number of 1s in a binary 0/1 dataset with a given number of patterns. Thus, *a tiling* is essentially a set of itemsets that together describe as many item occurrences as possible. Without loss of generality, we describe the task of sampling non-overlapping 2-tilings ($k = 2$). Let p_1 and p_2 denote the constituent patterns of a 2-tiling. The quality of a tiling is equal to its *area*, i.e., the number of 1s that it covers:

$$area\,(\{p_1, p_2\}) = (freq\,(p_1) \times |p_1| + freq\,(p_2) \times |p_2|)$$

The scaling constant for *area* is $C = \sum \mathcal{D}_{ti}$, i.e., the total number of 1s in the dataset. The tilt bound is $\hat{r} = \sum \mathcal{D}_{ti} / (2 \times (|\mathcal{D}| \times \theta) \times \lambda)$, where the denominator is the smallest possible area of a 2-tiling given the constraints.

7 Experiments

The experimental evaluation focuses on accuracy, scalability, and flexibility of the proposed sampler. Research questions are as follows:

Q1 *How close is the empirical sampling distribution to the target distribution?*

Q2 *How does* FLEXICS *compare to the specialized alternatives?*

Q3 *Does* FLEXICS *scale to large datasets?*

Q4 *How flexible is* FLEXICS*, i.e., can it be used for new pattern sampling tasks?*

The implementations of GFLEXICS and EFLEXICS are based on CP4IM[3] and a custom implementation of ECLAT respectively. Both are augmented with

[3] `https://dtai.cs.kuleuven.be/CP4IM`

	Constraints $\mathcal{C}$	Itemsets per dataset
F	$min\mathbf{F}req\,(\theta)$	$\sim 60\,000$
FCL	$min\mathbf{F}req\,(\theta) \wedge$ $\mathbf{C}losed \wedge min\mathbf{L}en\,(5)$	$\sim 20\,000$

Quality measure φ	$\hat{r}$
$uniform\ (\varphi \equiv 1)$	1
$purity$	2
$freq$	θ^{-1}

Table 3: We investigate sampling accuracy for all combinations of two constraint sets and three quality measures, yielding six experimental settings per dataset. See Section 4 for definitions.

a propagator for a system of XOR constraints based on the implementation of Gaussian elimination in the M4RI library[4]. All experiments were run on a Linux machine with an Intel Xeon CPU@3.2GHz and 32Gb of RAM.

Q1: Sampling accuracy We study the sampling accuracy of GFLEXICS in settings with tight constraints, which yield a relatively low number of solutions. This allows us to compute the exact statistical distance between the empirical distribution and the target distribution. We investigate settings with various quality measures and constraint sets as well as the effect of the tolerance parameter κ.

For each dataset, we construct six experimental settings by combining two constraint sets and three quality measures shown in Table 3. Table 5 shows dataset[5] statistics and values of θ. For each $\kappa \in \{0.1, 0.5, 0.9\}$, we requested 900 000 samples.

Let T denote the set of all itemsets that satisfy the constraints, E denote the multiset of all samples, and $\mathbf{1}_S$ its multiplicity function. For a given quality measure φ, target and empirical probabilities of sampling an itemset p are respectively defined as $P_T(p) = \varphi(p) / \sum_{p' \in T} \varphi(p')$ and $P_E(p) = \mathbf{1}_E(p) / |E|$. We use *Jensen-Shannon (JS) divergence* to quantify the statistical distance between P_T and P_E. Let $D_{KL}(P_1 \| P_2)$ denote the well-known *Kullback-Leibler divergence* between distributions P_1 and P_2. JS-divergence D_{JS} is defined as follows:

$$D_{JS}(P_T \| P_E) = 0.5 \times (D_{KL}(P_T \| P_M) + D_{KL}(P_Q \| P_M))$$
$$\text{where } P_M = 0.5 \times (P_T + P_E)$$

JS-divergence ranges from 0 to 1 and, unlike KL-divergence, does not require that $P_T(p) > 0 \Rightarrow P_E(p) > 0$, i.e., that each solution is sampled at least once, which does not always hold in sampling experiments. We compare D_{JS} attained with our sampler with that of the ideal sampler, which materializes all itemsets satisfying the constraints, computes their qualities, and uses these to sample directly from the target distribution.

[4] `https://bitbucket.org/malb/m4ri/`
[5] Source: `https://dtai.cs.kuleuven.be/CP4IM/datasets/`

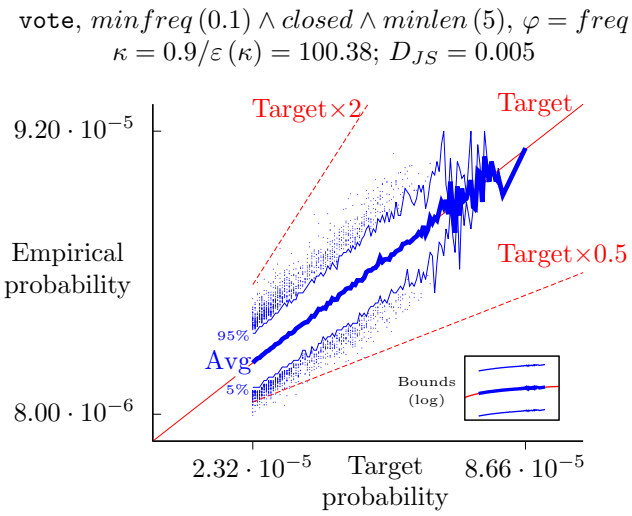

Figure 2: Empirical sampling frequencies of itemsets that share the same target probability, i.e., have the same quality. Average frequencies are close to target probabilities. 90% of frequencies are well within a factor 2 from the target, which is considerably lower than the theoretical factor of 100.38. (The dots show the tails of the empirical probability distribution for a given target probability. The box shows theoretical bounds and empirical frequencies on the logscale).

A characteristic experiment in detail Our experiments show that results are consistent across various datasets. Therefore, we first study the results on the `vote` dataset in detail. Table 4 shows that the theoretical error tolerance parameter κ has no considerable effect on practical performance of the algorithm, except for runtime, which we evaluate in subsequent experiments. One possible explanation is the high quality of the output of the estimation phase, which thus alleviates theoretical risks that have to be accounted for in the general case (see below for a numerical characterization). Hence, in the following experiments we use $\kappa = 0.9$ unless noted otherwise.

JS-divergences for different quality measures and constraint sets are impressively low, equivalent to the highest possible sampling accuracy attainable with the ideal sampler. Figure 2 illustrates this for $minfreq\,(\theta) \wedge closed \wedge minlen\,(5)$, $\varphi = freq$, and $\kappa = 0.9$ ($D_{JS} = 0.005$): the sampling frequency of an average itemset is close to the target probability. For at least 90% of patterns, the sampling error does not exceed a factor of 2.

Table 5 shows that similar conclusions hold for several other datasets. Over all 24 experimental settings, the error of the estimation of the total weight of all solutions, which is used to derive the number of XOR constraints for the sampling phase, never exceeds 10%, whereas the bounds assume the error of 45 to 80%. This helps explain why practical errors are considerably lower than theoretical bounds.

Table 4: Sampling accuracy of FLEXICS (here GFLEXICS) is consistently high across quality measures, constraint sets ($min\mathbf{F}req\,(0.1)$ vs. $min\mathbf{F}req\,(0.1) \wedge \mathbf{C}losed \wedge min\mathbf{L}en\,(5)$), and error tolerance κ. JS-divergence is impressively low, equivalent to that of the ideal sampler.

	`vote` dataset, JS-divergence from target					
	Uniform ($\hat{r} = 1$)		Purity ($\hat{r} = 2$)		Frequency ($\hat{r} = 10$)	
κ	F	FCL	F	FCL	F	FCL
0.9	0.009	0.005	0.009	0.005	0.009	0.005
0.5	0.009	0.005	0.009	0.005	0.009	0.005
0.1	0.009	0.005	0.009	0.005	0.009	0.005
Ideal sampler	0.009	0.005	0.009	0.005	0.009	0.005

Table 5: Dataset-specific statistics and parameter settings and results of sampling accuracy experiments. Even with high error tolerance $\kappa = 0.9$, JS-divergence is consistently low across datasets, quality measures, and constraint sets.

					JS-divergence, $\kappa = 0.9$					
					Uniform		Purity		Frequency	
	$\lvert\mathcal{D}\rvert$	$\lvert\mathcal{I}\rvert$	Density	θ	F	FCL	F	FCL	F	FCL
`heart`	296	95	47%	0.44 (130)	0.009	0.002	0.010	0.002	0.010	0.002
`hepatitis`	137	68	50%	0.40 (55)	0.009	0.003	0.009	0.003	0.009	0.003
`primary`	336	31	48%	0.09 (30)	0.012	0.006	0.012	0.006	0.012	0.006
`vote`	435	48	33%	0.10 (44)	0.009	0.005	0.009	0.005	0.009	0.005

Q2: Comparison with alternative pattern samplers We compare FLEXICS with ACFI and TS, alternative samplers described in Section 3, however only in the settings that they are tailored for[6]. ACFI only supports the setting with a single $minfreq\,(\theta)$ constraint and $\varphi = uniform$. It is run with a burn-in of 100,000 steps and uses a built-in heuristic to determine the number of steps between consecutive samples. TS is evaluated in the setting with $\varphi = freq$ and both constraint sets from the previous experiments. It samples from two of the distributions it supports, $freq$ and $freq^4$; samples that do not satisfy the constraints are rejected. Both samplers are requested to generate 900,000 samples and are allowed to run up to 24 hours. Datasets and parameters are identical to the previous experiments.

Table 6 shows the accuracy of various samplers. The performance of FLEXICS is on par with specialized samplers. That is, in uniform frequent itemset

[6]The code was provided by their respective authors. We also obtained the "unmaintained" code for the *uniform LRW* sampler (personal communication), but were unable to make it run on our machines. We were unable to obtain the code for the *FCA* sampler.

Table 6: The accuracy of FLEXICS (here GFLEXICS) is consistent across settings. In uniform frequent itemset sampling, performance of FLEXICS as well as of ACFI is equivalent to that of the ideal sampler (not shown). In frequency-weighted sampling, it is comparable to the exact two-step sampler (TS $\sim freq$) with rejection. However, the latter suffers from low acceptance rates, which, for settings marked with '−', is not improved by increasing its bias (TS $\sim freq^4$).

	JS-divergence (for TS, acceptance rate)							
	Uniform		Frequency					
	F		F			FCL		
	GF	ACFI	GF	TS $\sim freq$	TS $\sim freq^4$	GF	TS $\sim freq$	TS $\sim freq^4$
`heart`	0.009	0.009	0.010	− (0%)	− (0%)	0.002	− (0%)	− (0%)
`hepatitis`	0.009	0.009	0.009	− (≪ 0.1‰)	− (0.9%)	0.003	− (≪ 0.1‰)	− (0.3%)
`primary`	0.012	0.013	0.012	0.012 (30%)	0.392 (39%)	0.006	0.006 (16%)	0.273 (39%)
`vote`	0.009	0.010	0.009	0.010 (11%)	0.195 (91%)	0.005	0.005 (7%)	0.136 (48%)

sampling, the accuracy of both FLEXICS and ACFI is equivalent to that of the ideal sampler and can therefore not be improved. When sampling proportional to frequency, it is equivalent to the accuracy of the exact two-step sampler TS $\sim freq$. However, the latter does not directly take constraints into account, which poses considerable problems on the `heart` and `hepatitis` datasets. For example, for the `heart` dataset, TS fails to generate a single accepted sample, despite generating 2 billion unconstrained candidates. This issue is not solved by increasing the bias towards more frequent itemsets by sampling proportional to $freq^4$. Furthermore, this would substantially decrease accuracy, as seen in `primary` and `vote`.

Table 7 shows the runtimes for frequent itemset sampling (i.e., only the $minfreq$ constraint). In most settings, EFLEXICS provides runtime benefits over GFLEXICS. Regardless, both are outperformed by specialized samplers. In uniform sampling, a slower variant is from 6 to 22 milliseconds slower than ACFI. In frequency-weighted sampling, FLEXICS is considerably faster in the settings with tighter constraints, where the two-step sampler is slow to generate accepted samples. This illustrates the overhead as well as the benefits of the flexibility of the proposed approach. Furthermore, in these settings, there are at most 63 000 patterns, which is too low to suggest the need for pattern sampling (recall that the primary goal of these experiments was to evaluate and compare sampling accuracy) and does not allow for the overhead amortization. We therefore tackle settings with a substantially larger number of patterns in the following experiments.

Q3: Scalability To study scalability of the proposed sampler, we compare its runtime costs with those required to construct an ideal sampler with LCM[7], an efficient frequent itemset miner [31]. To this end, we estimate the costs

[7] `http://research.nii.ac.jp/~uno/codes.htm`, ver. 3

Table 7: Runtime in milliseconds required to sample a frequent itemset, including pre-processing, i.e., estimation or burn-in, amortized over 1000 samples. Both variants of FLEXICS are suitable for anytime exploration, although slower than the specialized samplers. The two-step sampler is the fastest in the task it is tailored for, but fails in the settings with tighter constraints. EFLEXICS provides runtime benefits compared to GFLEXICS.

	$\varphi = uniform, \mathcal{C} = \mathrm{F}$			$\varphi = freq, \mathcal{C} = \mathrm{F}$		
	GFLEXICS	EFLEXICS	ACFI	GFLEXICS	EFLEXICS	TS $\sim freq$
`heart`	54	38	32	65	44	–
`hepatitis`	18	28	12	23	34	3638
`primary`	9	10	4	26	26	0.3
`vote`	22	19	8	40	28	1.1

of completing the following scenario: pre-processing (estimation or counting), followed by sampling 100 itemsets in two batches of 50. Sampling accuracy of FLEXICS is evaluated indirectly, via the results of the estimation phase.

We use dataset `accidents`[8] (469 items, 340183 transactions) and $\theta = 0.009$ (3000 transactions), which results in a staggering number of 5.37 billion frequent itemsets. We run WEIGHTGEN with values of $\kappa \in \{0.1, 0.5, 0.9\}$. (Note that the estimation phase is identical for all three cases.) The baseline sampler is constructed as follows. LCM is first run in counting mode, which only returns the total number of itemsets. Then, for each batch, 50 random line numbers are drawn, and the corresponding itemsets are printed while LCM is enumerating the solutions[9]. The latter phase is implemented with the standard Unix utility '`awk`'.

Figure 3 illustrates the results. The counting mode of LCM is roughly 4.5 minutes faster than the estimation phase of EFLEXICS. Generating samples from the output of LCM, on the other hand, is considerably slower: it takes approximately 35s to sample one itemset, whereas EFLEXICS takes from 10s to 27s per sample, depending on error tolerance κ. As a result, EFLEXICS samples two batches faster than LCM regardless of its parameter values. Moreover, with $\kappa = 0.9$ it samples all 100 itemsets even before the first batch is returned by LCM.

Thus, the proposed sampler outperforms a sampler derived from an efficient itemset miner, even though the experimental setup favors the latter. First, non-uniform weighted sampling would require more advanced computations with itemsets, which would increase the costs of both counting and sampling with LCM. Second, EFLEXICS could also benefit from the exact count obtained by LCM and start sampling after 1.5 minutes. Third, the individual itemsets sampled from the output of an algorithm based on deterministic search are not

[8] Source: `http://fimi.ua.ac.be/data/`

[9] Storing all itemsets on disk provides no benefits: it increases the mining runtime to 23 minutes and results in a file of 215Gb; simply counting its lines with '`wc -l`' takes 25 minutes.

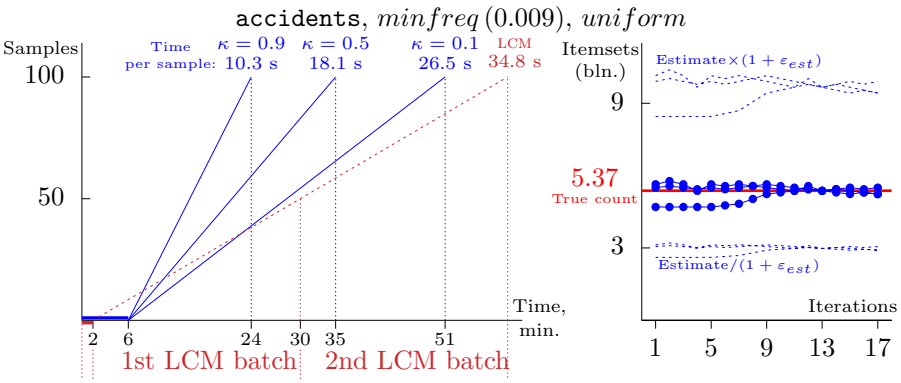

a) Sampling runtime comparison b) Estimation accuracy

Figure 3: a) EFlexics generates two batches of 50 samples faster than a sampler derived from lcm, regardless of error tolerance. b) EFlexics with the *uniform* quality converges to a high-quality estimate of the total number of itemsets in a small number of iterations (three different random seeds shown). Practical error of the estimation phase is substantially lower than theoretical bounds, which indirectly signals high sampling accuracy.

exchangeable. Figure 4 illustrates this: due to lcm's search order, certain items only occur at the beginning of batches, while for EFlexics, the order within a batch is random.

Accuracy of Flexics in this scenario can be evaluated indirectly, by comparing the estimate of the total number of itemsets obtained at the estimation phase with the actual number. The error tolerance of the estimation phase is $\varepsilon_{est} = 0.8$ (see Appendix A for details). Figure 3b demonstrates that, in practice, the error is substantially lower than the theoretical bound. Furthermore, 3 to 9 iterations suffice to obtain an accurate estimate. Similar to previous experiments, accurate input from the estimation phase alleviates theoretical risks and is expected to enable accurate sampling.

Q4: Pattern set sampling In order to demonstrate the flexibility of our approach and the promised benefits of weighted constrained pattern sampling, i.e., 1) diversity and quality of results, 2) utility of constraints, and 3) the potential for anytime exploration, we here address the problem of sampling non-overlapping 2-tilings as introduced in Section 6. We re-use the implementation of GFlexics from the itemset sampling experiments, only modifying the declarative specification of the CSP. Constraints on individual patterns are identical to previous experiments: $minfreq(\theta) \wedge closed \wedge minlen(5)$.

Table 8 shows parameters and runtimes for sampling 2-tilings proportional to *area*. In all cases, it takes from 3 to 5 seconds to sample one 2-tiling, which is suitable for pattern-based data exploration, where tilings are inspected by a human analyst. For several settings, the estimation phase overhead slightly exceeds the runtime of complete enumeration of all solutions. However, for the

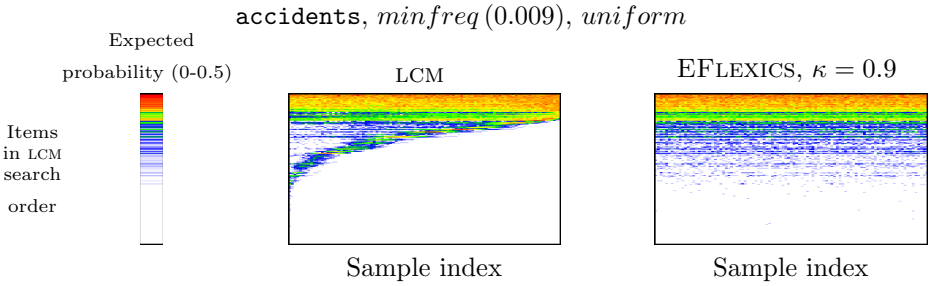

Figure 4: The probability of observing a given item at a certain position in a batch by EFLEXICS is close to the expected probability of observing this item in a random itemset, which indicates high sampling accuracy. The samples by the LCM-based sampler are not exchangeable, i.e., certain items are under- or oversampled at certain positions in a batch, depending on their position in LCM's search order.

Table 8: Time required to sample a 2-tiling is approximately 4s, which is suitable for anytime exploration. Runtime benefits of the sampling procedure are the largest for the settings with the largest tiling counts (`primary` and `vote`).

					Sampling with GFLEXICS	
	θ	Tilt bound $\hat{r}$	Tilings, mln.	Enumeration, min	Estimation, min	Per sample, s
`heart`	0.30	13.3	2.2	0.5	2.2	3.8
`hepatitis`	0.26	12.4	7.2	0.9	1.8	3.4
`primary`	0.03	50.3	24.9	3.0	2.5	4.3
`vote`	0.10	15.3	170.1	20.8	2.3	4.8

settings with a large number of pattern sets, which are arguably the primary target of pattern samplers, the opposite is true. For example, in the `vote` experiment with 170 mln. tilings, the runtime of the estimation phase only amounts to 11% of the runtime of complete enumeration, which demonstrates the benefits of the proposed approach.

Figure 5 shows six random 2-tilings sampled from the `vote` dataset. Constraints ensure that the individual tiles comprising each 2-tiling do not overlap, simplifying interpretation. Moreover, the *set* of tilings is diverse, i.e., the tilings are dissimilar to each other. They cover different regions in the data, revealing alternative structural regularities.

The right part of Figure 5 shows the *area* distribution of all 2-tilings that satisfy the constraints, obtained by complete enumeration. Qualities of 5 out of 6 tilings fall in the dense region between the 25th and 75th percentile, indicating high sampling accuracy. This is completely expected from the problem statement. In practice, pattern quality measures, like *area*, are only an approx-

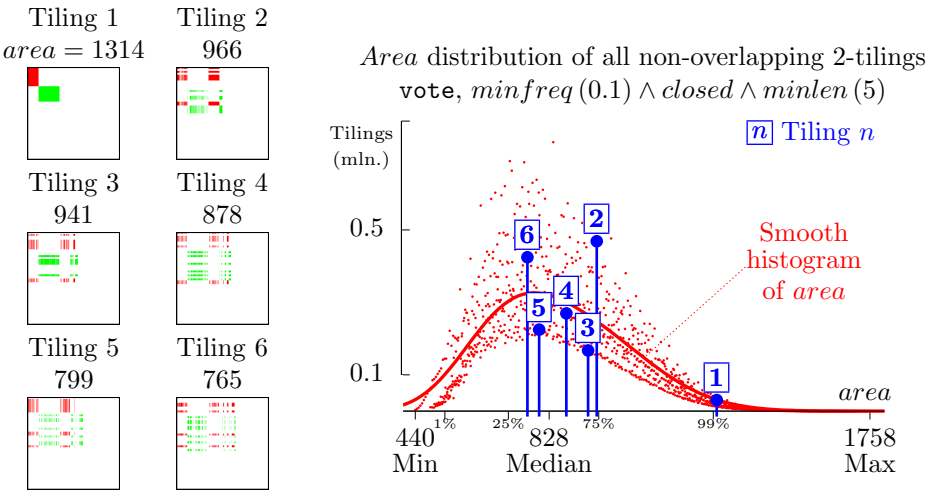

Figure 5: **Left**: Six 2-tilings sampled consecutively from the `vote` dataset. The tilings are diverse, i.e., cover different regions in the data, a property essential for pattern-based data exploration. (Note that while the sampled tilings are fair random draws, the images are not random: the tilings were sorted by *area* descending, and items and transactions were re-arranged so that the cells covered by tilings with larger *area* are as close to each other as possible.) **Right**: Qualities (*area*) of the samples, indicated by vertical bars, tend towards a dense region between the 25th and the 75th percentile.

imation of application-specific pattern interestingness, thus diversity of results is a desirable characteristic of a pattern sampler as long as the quality of individual patterns is sufficiently high. To sample patterns from the right tail (i.e., with exceptionally high qualities) more frequently, the sampling task could be changed, e.g., either by choosing another sampling distribution or by enforcing constraints on *area*.

8 Conclusion

We proposed FLEXICS, a flexible pattern sampler with theoretical guarantees regarding sampling accuracy. We leveraged the perspective on pattern mining as a CSP and developed the first pattern sampling algorithm that builds upon the latest advances in sampling solutions in SAT.

Experiments demonstrate that FLEXICS delivers the promised benefits: 1) it is flexible in that it supports a wide range of pattern constraints and sampling distributions in itemset mining as well as the novel pattern set sampling task; 2) it is anytime in that the time it takes to generate random patterns is suitable for online data exploration, including the settings with large datasets or large solution spaces; 3) by virtue of high sampling accuracy in all supported

settings, sampled patterns are diverse, i.e., originate from different regions in the solution space. The theoretical guarantees ensure that these conclusions extend reliably beyond the studied settings. Furthermore, practical accuracy is substantially higher than theory guarantees. Finally, our approach provides a principled method to convert a pattern miner into a pattern sampler, which we accomplished with CP4IM and ECLAT.

Directions for future work include extensions to richer pattern languages and relaxing assumptions regarding sampling distributions. The two settings that we have studied in this paper can be compactly encoded with binary variables. While in principle, richer pattern languages, e.g., sequences or graphs, could be modeled using this framework, the number of variables would rise sharply, which would negatively affect performance. A possible solution would use generalized XOR constraints with integer variables [28].

The bounded tilt assumption for sampling distributions could be tackled with divide-and-conquer approaches [16, Section 6] or alternative estimation techniques [32]. This would require the capacity to handle constraints on bounds of a pattern quality measure efficiently, which is possible for certain quality measures [33, 34].

References

[1] Rakesh Agrawal, Heikki Mannila, Ramakrishnan Srikant, Hannu Toivonen, and A. Inkeri Verkamo. *Advances in Knowledge Discovery and Data Mining*, chapter Fast Discovery of Association Rules, pages 307–328. 1996.

[2] Charu C Aggarwal and Jiawei Han, editors. *Frequent Pattern Mining*. Springer International Publishing, 2014.

[3] Toon Calders, Christophe Rigotti, and Jean-François Boulicaut. A survey on condensed representations for frequent sets. In Jean-François Boulicaut, Luc De Raedt, and Heikki Mannila, editors, *Constraint-Based Mining and Inductive Databases*, pages 64–80. Springer Berlin Heidelberg, 2006.

[4] Albrecht Zimmermann and Siegfried Nijssen. Supervised pattern mining and applications to classification. In C. Charu Aggarwal and Jiawei Han, editors, *Frequent Pattern Mining*, chapter 17, pages 425–442. Springer International Publishing, 2014.

[5] Siegfried Nijssen and Albrecht Zimmermann. Constraint-based pattern mining. In C. Charu Aggarwal and Jiawei Han, editors, *Frequent Pattern Mining*, chapter 7, pages 147–163. Springer International Publishing, 2014.

[6] Björn Bringmann, Siegfried Nijssen, Nikolaj Tatti, Jilles Vreeken, and Albrecht Zimmermann. Mining sets of patterns. *Tutorial at ECML/PKDD*, 2010.

[7] Mario Boley and Henrik Grosskreutz. Approximating the number of frequent sets in dense data. *Knowledge and information systems*, 21(1):65–89, 2009.

[8] Mohammad Al Hasan and Mohammed Zaki. Output space sampling for graph patterns. In *Proceedings of VLDB*, pages 730–741, 2009.

[9] Mario Boley, Thomas Gärtner, and Henrik Grosskreutz. Formal Concept Sampling for Counting and Threshold-Free Local Pattern Mining. In *Proceedings of SDM*, pages 177–188, 2010.

[10] Mario Boley, Claudio Lucchese, Daniel Paurat, and Thomas Gärtner. Direct Local Pattern Sampling by Efficient Two-Step Random Procedures. In *Proceedings of KDD*, pages 582–590, 2011.

[11] Mario Boley, Sandy Moens, and Thomas Gärtner. Linear space direct pattern sampling using coupling from the past. In *Proceedings of KDD*, pages 69–77, 2012.

[12] Carla P Gomes, Ashish Sabharwal, and Bart Selman. Near-uniform sampling of combinatorial spaces using XOR constraints. In *Proceedings of NIPS*, pages 481–488, 2006.

[13] Supratik Chakraborty, Kuldeep S Meel, and Moshe Y Vardi. A scalable and nearly uniform generator of sat witnesses. In *Proceedings of CAV*, pages 608–623, 2013.

[14] Stefano Ermon, Carla P Gomes, Ashish Sabharwal, and Bart Selman. Embed and Project: Discrete Sampling with Universal Hashing. In *Proceedings of NIPS*, pages 2085–2093, 2013.

[15] Kuldeep Meel, Moshe Vardi, Supratik Chakraborty, Daniel Fremont, Sanjit Seshia, Dror Fried, Alexander Ivrii, and Sharad Malik. Constrained Sampling and Counting: Universal Hashing Meets SAT Solving. In *Proceedings of Beyond NP AAAI Workshop*, 2016.

[16] Supratik Chakraborty, Daniel J Fremont, Kuldeep S Meel, and Moshe Y Vardi. Distribution-Aware Sampling and Weighted Model Counting for SAT. In *Proceedings of AAAI*, pages 1722–1730, 2014.

[17] Tias Guns, Siegfried Nijssen, and Luc De Raedt. Itemset mining: A constraint programming perspective. *Artificial Intelligence*, 175(12-13):1951–1983, aug 2011.

[18] Mohammed J. Zaki, Srinivasan Parthasarathy, Mitsunori Ogihara, and Wei Li. New algorithms for fast discovery of association rules. In *Proceedings of KDD*, pages 283–296, 1997.

[19] Jian Pei and Jiawei Han. Can we push more constraints into frequent pattern mining? In *Proceedings of KDD*, pages 350–354, 2000.

[20] Luc De Raedt and Albrecht Zimmermann. Constraint-based pattern set mining. In *Proceedings of SDM*, pages 237–248, 2007.

[21] Arno Knobbe, Bruno Crémilleux, Johannes Fürnkranz, and Martin Scholz. From local patterns to global models: the LeGo approach to data mining. In *Proceedings of LeGo ECML PKDD Workshop*, pages 1–16, 2008.

[22] Cristian Bucilă, Johannes Gehrke, Daniel Kifer, and Walker White. Dualminer: A dual-pruning algorithm for itemsets with constraints. *Data Mining and Knowledge Discovery*, 7(3):241–272, 2003.

[23] Francesco Bonchi, Fosca Giannotti, Claudio Lucchese, Salvatore Orlando, Raffaele Perego, and Roberto Trasarti. A constraint-based querying system for exploratory pattern discovery. *Information Systems*, 34(1):3–27, 2009.

[24] Mehdi Khiari, Patrice Boizumault, and Bruno Crémilleux. Constraint programming for mining n-ary patterns. In *Proceedings of CP*, pages 552–567, 2010.

[25] Tias Guns, Siegfried Nijssen, and Luc De Raedt. k-pattern set mining under constraints. *Knowledge and Data Engineering, IEEE Transactions on*, 25(2):402–418, 2013.

[26] A Kemmar, W Ugarte, S Loudni, T Charnois, Y Lebbah, P Boizumault, and B Crémilleux. Mining relevant sequence patterns with cp-based framework. In *Proceedings of ICTAI*, pages 552–559, 2014.

[27] Sergey Paramonov, Matthijs van Leeuwen, Marc Denecker, and Luc De Raedt. An exercise in declarative modeling for relational query mining. In *Proceedings of ILP*, 2015.

[28] Carla P Gomes, Willem-jan van Hoeve, Ashish Sabharwal, and Bart Selman. Counting CSP Solutions Using Generalized XOR Constraints. In *Proceedings of AAAI*, pages 204–209, 2007.

[29] Mate Soos. Enhanced Gaussian Elimination in DPLL-based SAT Solvers. In *Proceedings of Pragmatics of SAT Workshop*, 2010.

[30] Floris Geerts, Bart Goethals, and T Mielikäinen. Tiling databases. In *Proceedings of Discovery Science*, pages 278–289, 2004.

[31] Takeaki Uno, Masashi Kiyomi, and Hiroki Arimura. LCM ver. 3: Collaboration of array, bitmap and prefix tree for frequent itemset mining. In *Proceedings of OSDM Workshop*, pages 77–86, 2005.

[32] Stefano Ermon, Carla P. Gomes, Ashish Sabharwal, and Bart Selman. Taming the curse of dimensionality: Discrete integration by hashing and optimization. In *Proceedings of ICML*, pages 334–342, 2013.

[33] Siegfried Nijssen, Tias Guns, and Luc De Raedt. Correlated Itemset Mining in ROC Space: A Constraint Programming Approach. In *Proceedings of KDD*, pages 647–655, 2009.

[34] Florian Lemmerich, Martin Atzmueller, and Frank Puppe. Fast exhaustive subgroup discovery with numerical target concepts. *Data Mining and Knowledge Discovery*, 30(3):711–762, 2016.

[35] Supratik Chakraborty, Daniel J. Fremont, Kuldeep S. Meel, Sanjit A. Seshia, and Moshe Y. Vardi. On parallel scalable uniform SAT witness generation. In *Proceedings of TACAS*, volume 9035, pages 304–319, 2015.

A WeightGen

In this section, we present an extended technical description of the WeightGen algorithm, which closely follows Sections 3 and 4 in [16], whereas the pseudocode in Algorithm 2 is structured similarly to that of UniGen2, a close cousin of WeightGen [35]. Lines 1-3 correspond to the estimation phase and Lines 4-8 correspond to the sampling phase. SolveBounded stands for the bounded enumeration oracle.

The parameters of the estimation phase are fixed to particular theoretically motivated values. $pivot_{est}$ denotes the maximal weight of a cell at the estimation phase; $pivot_{est} = 46$ corresponds to estimation error tolerance $\varepsilon_{est} = 0.8$ (Line 10). If the total weight of solutions in a given cell exceeds $pivot_{est}$, a new random XOR constraint is added in order to eliminate a number of solutions. Repeating the process for a number of iterations increases the confidence of the estimate, e.g. 17 iterations result in $1-\delta_{est} = 0.8$ (Line 1). Note that Estimate essentially estimates the total weight of *all* solutions, from which N_{XOR}, the initial number of XOR constraints for the sampling phase, is derived (Line 4).

A similar procedure is employed at the sampling phase. It starts with N_{XOR} constraints and adds at most *three* extra constraints. The user-chosen error tolerance parameter κ determines the range $[loThresh,\ hiThresh]$, within which the total weight of a suitable cell should lie (Line 5). For example, $\kappa = 0.9$ corresponds to range [6, 18]. If a suitable cell can be obtained, a solution is sampled exactly from all solutions in the cell; otherwise, no sample is returned. Requiring the total cell weight to exceed a particular value ensures the lower bound on the sampling accuracy.

The preceding presentation makes two simplifying assumptions: (1) all weights lie in $[1/r,\ 1]$; (2) adding XOR constraints never results in unsatisfiable subproblems (empty cells). The former is relaxed by multiplying pivots by $\hat{w}_{max} = \hat{w}_{min} \times \hat{r} < 1$, where $\hat{w}_{min}$ is the smallest weight observed so far. The latter is solved by simply restarting an iteration with a newly generated set of constraints. See [16] for the full explanation, including the precise formulae to compute all parameters.

Implementation details Following suggestions of [35], we implement *leapfrogging*, a technique that improves the performance of the umbrella sampling procedure

and thus benefits both GFlexics and EFlexics. First, after three iterations of the estimation phase, we initialize the following iterations with a number of XOR constraints that is equal to the smallest number returned in the previous iterations (rather than with zero XORs). Second, in the sampling phase, we start with one XOR constraint more than the number suggested by theory. If the cell is too small, we remove one constraint; if it is too large, we proceed adding (at most two) constraints. Both modifications are based on the observation that theoretical parameter values address hypothetical corner cases that rarely occur in practice. Finally, we only run the estimation phase until the initial number of XOR constraints, which only depends on the median of total weight estimates, converges. For example, if the estimation phase is supposed to run for 17 iterations, the convergence can happen as early as after 9 iterations.

Algorithm 2 WeightGen [16]

Input: Boolean formula F, weight w, tilt bound $\hat{r}$, sampling error tolerance parameter κ

Assumes: $w(\cdot) \in [1/\hat{r},\ 1]$, bounded enumeration algorithm SolveBounded

for 17 iterations **do** ▷ Corresponds to $\delta_{est} = 0.2$
  $WeightEstimates \overset{Add}{\leftarrow}$ Estimate($\emptyset$)
$TotalWeight =$ Median($WeightEstimates$)
$N_{XOR} = \mathcal{O}\left(\log_2 TotalWeight / \left(1 + \kappa^{-1}\right)\right)$
$loThresh \propto (1+\kappa)/\kappa^2$, $hiThresh \propto (1+\kappa)^3/\kappa^2$
for $N_{samples}$ times **do**
  $InitXORs = \{$RandomXOR() $\times N_{XOR}$ times$\}$
  Generate(κ, $[loThresh,\ hiThresh]$, $InitXORs$, 3)

function Estimate($XORs$)
▷ Returns an estimate of the total weight of all solutions
  $pivot_{est} = 46$ ▷ Corresponds to $\varepsilon_{est} = 0.8$
  $Sols \leftarrow$ SolveBounded(F, $XORs$, $pivot_{est}$)
  $CellWeight \leftarrow \sum_{s \in Sols} w(s)$
  if $CellWeight \leq pivot_{est}$ **then** ▷ Cell of the "right" size
    return $CellWeight \times 2^{|XORs|}$
  else ▷ Shrink cell by adding XOR constraint
    Estimate($XORs \cup$ RandomXOR())

function Generate(κ, $[lT,\ hT]$, $XORs$, i)
▷ Returns a random solution of F
  $Sols \leftarrow$ SolveBounded(F, $XORs$, hT)
  $CellWeight \leftarrow \sum_{s \in Sols} w(s)$
  if $CellWeight \in [lT,\ hT]$ **then** ▷ Cell of the "right" size
    return SampleExactly($Sols$, w)
  else if $CellWeight > lT \wedge i > 0$ **then** ▷ Cell is too large
    Generate(κ, $[lT,\ hT]$, $XORs \cup$ RandomXOR(), $i-1$)
  else ▷ Cell is too small
    return $\perp$